

Interpreting presuppositions in the scope of quantifiers: *Every* vs. *at least one*¹

Ava CREEMERS — *University of Pennsylvania*

Jérémy ZEHR — *University of Pennsylvania*

Florian SCHWARZ — *University of Pennsylvania*

Abstract. This paper experimentally investigates presupposition projection from the scope of the quantifiers *every* and *at least one*, as triggered by the factive verb *be aware* and the iterative adverb *again*. The first issue we are concerned with is whether presuppositions project universally or existentially from quantified sentences. Different theoretical accounts endorse opposing views here (e.g., Heim, 1983; Geurts, 1999; Beaver, 2001; Schlenker, 2008, 2009; Fox, 2012), while recent experimental work (Chemla, 2009; Tiemann, 2014) suggests that the force of the projected presupposition varies by quantifier. The second issue we look at is how the descriptively observed readings arise—in particular, as a direct result output from the projection mechanism, or via additional, independent mechanisms such as domain restriction (e.g., Geurts and van Tiel, 2016): if the domain of the quantifier is restricted, this can yield what looks like non-universal inferences in light of the overall, unrestricted domain, even if the projection mechanism itself yields a universal presupposition. Finally, we test whether the presupposed content also forms part of the entailed content, at least for certain triggers (Sudo, 2012; Klinedinst, 2016; Zehr and Schwarz, 2016). Our results yield clearly different patterns for *every* and *at least one*, with *every* giving rise to universal presuppositions, which, to a very limited extent, can be weakened by domain restriction, and *at least one* overwhelmingly giving rise to non-universal presuppositions. Our results also indicate the availability of presupposition-less readings for both triggers in the task at hand, apparently more prevalent than domain restriction. Thereby, we present novel evidence that helps to pinpoint which of the theoretical options can be substantiated experimentally.

Keywords: Presupposition projection, quantifiers, domain restriction, entailment.

1. Introduction

One of the core properties of presuppositions is that they generally project out of a variety of embedding environments which cancel entailed content. For example, (1a-c), with the factive verb *be aware* as a presupposition trigger, uniformly presuppose that *the alien is blue*, despite being embedded under negation or in a question, both of which cancel the entailed content of the embedded material (Karttunen, 1973).

- (1) $\left\{ \begin{array}{l} \text{a. } \underline{\text{The alien}} \text{ is } \mathbf{aware} \text{ that } \underline{\text{he is blue}} \\ \text{b. } \underline{\text{The alien}} \text{ is not } \mathbf{aware} \text{ that } \underline{\text{he is blue}} \\ \text{c. Is } \underline{\text{the alien}} \mathbf{aware} \text{ that } \underline{\text{he is blue?}} \end{array} \right\} \quad \rightsquigarrow \text{The alien is blue}$

¹We gratefully acknowledge support from NSF-grant BCS-1349009 to Florian Schwarz. We thank the audience at the workshop *Theoretical and Experimental Approaches to Presuppositions* in Genoa, the audience at Sinn und Bedeutung 22, and the members at the lab meeting at UPenn for their comments.

However, when the presupposition trigger appears in the scope of a quantifier, there are opposing views as to whether presuppositions project universally or existentially (cf. Chemla, 2009). Some argue that quantified sentences as in (2) give rise to a universal presupposition (Heim, 1983; Schlenker, 2008, 2009), i.e., that every alien is in fact blue (3a). In contrast, others argue that semantic presuppositions of all quantified sentences are existential (Beaver, 2001, 1994); in our example that there exists at least one alien that is blue (3b).

- (2) a. Every alien is aware that he is blue
 b. At least one alien is aware that he is blue
- (3) a. $\rightsquigarrow^? \forall x \in D_{alien} blue(x)$ (\forall presupposition)
 b. $\rightsquigarrow^? \exists x \in D_{alien} blue(x)$ (\exists presupposition)

Recent experimental work (Chemla, 2009; Tiemann, 2014) suggests that the force of the projected presupposition varies by quantifier. Chemla (2009) uses an inference paradigm to investigate projection from the scope of quantifiers in French. The results indicate a significant difference in the availability of universal presuppositions, depending on the quantifier used: he finds evidence for universal projection from the scope of the universal quantifier *each*, but not from the scope of existential quantifiers (*less than 3*, *more than 3*, *exactly 3*). The perhaps most contested case is the negative quantifier *no*: Chemla argues that his data support universal projection, but more recent work by Zehr et al. (2016) provides evidence for (at least the possibility of) existential projection. Similarly, Tiemann (2014), in an eye-tracking experiment, shows that reading measures differ significantly depending on whether a universal or an existential quantifier is used. Together, these studies suggest that presuppositions do not uniformly project universally or existentially—rather, the projection behavior changes with the quantifier.

In this paper, we report data on how presuppositions triggered by *be aware* and *again* project from the scope of the quantifiers *every* (2a) and *at least one* (2b). Moreover, we test whether the descriptively observed existential projection readings are derived directly via the projection mechanism, or whether they are derived from the output of the projection mechanism through other processes. Candidates that could be at play include (implicit) Domain Restriction (Geurts and van Tiel, 2016) and the inclusion of presupposed content at the assertive level (Sudo, 2012). Our results confirm previous results in that the quantifiers *every* and *at least one* pattern differently, with *every* giving rise to universal readings of the presupposition, which only can be weakened by domain restriction to a limited extent, and *at least one* giving rise to non-universal readings of the presupposed content. This provides support for theories that tie different projection behavior to the nature of the quantifier at play, rather than treating all quantifiers as having uniform projection behavior. Furthermore, our results indicate that, at least within our task paradigm, presupposition-less readings are available for both triggers, and this option seems to be more prevalent than domain restriction.

The paper is structured as follows. Section 2 presents the background on additional mechanisms that force what seems like an existential presupposition: domain restriction and (non-)entailment. Section 3 presents the methods of our experiment, and Section 4 gives the results. Section 5 discusses the theoretical implications of the results, and Section 6 concludes.

2. Background

This section discusses two important factors that can affect whether presupposed content under quantifiers give rise to universal or existential inferences: implicit Domain Restriction (Section 2.1) and (non-)entailment of the presupposition (Section 2.2). Section 2.3 proceeds to lay out the rationale for the experiments in the present study.

2.1. (Implicit) Domain Restriction

Presuppositions triggered within the scope of a quantifier may restrict the domain of individuals considered in evaluating the quantificational claim, a mechanism known as (implicit) Domain Restriction (henceforth DR). In the presence of DR, a presupposition may be universally satisfied relative to the restricted domain, while appearing non-universal in light of the unrestricted domain. Let us illustrate with the sentence in (2a). A universal presupposition gives rise to the notion that all aliens are in fact blue. This would necessarily be incompatible with a case in which there are non-blue aliens, as in Figure 1.




Figure 1: Illustration of Domain Restriction yielding seemingly non-universal inferences

However, with implicit DR, the quantifier's domain could be restricted to those aliens that are in fact blue (i.e., the five blue aliens on the right, but crucially not the green aliens on the left in Figure 1). If such a DR is available, it should be possible to judge the sentence compatible with situations with non-blue aliens, like that in Figure 1. Therefore, if DR can come into play in the relevant sentences, it becomes non-trivial to determine whether or not the projection mechanism indeed gives rise to a universal presupposition, as what looks like a non-universal reading relative to the full domain in fact could result from interpreting the relevant presupposition universally relative to a restricted domain (Schlenker, 2008; Rothschild, 2011; Sudo, 2012). Schematically, our sentence in (2a) can effectively be understood as in (4), resulting in the inference pattern represented in (5).

(4) Every alien [**that is blue**] is aware that he is blue.

(5) $\forall + \text{DR: } \forall x \in D_{\text{blue-alien}} [\text{blue}(x)] \approx \exists: \exists x \in D [\text{blue}(x)]$

In a recent experimental study, Geurts and van Tiel (2016) investigate the effects of presuppositions on DR. Specifically, they investigate the possibility of restricting the domain of universal quantification to those individuals that satisfy the presupposition of the scope of the quantifier. In a series of truth value judgment tasks, they paired simple geometrical figures (as illustrated

in Figure 2) with quantified sentences of the form *Each of these circles has the same color as the square to which it is connected*, in which “the square to which...” is the critical presupposition trigger. Crucially, the results show that sentences of this type are judged true 87% of the time when paired with a picture in which only four out of five circles were connected to a square (and have the same color as the square). The authors argue that this substantial amount of acceptances shows that the domain of quantification can be restricted by contextual factors. (In addition, their results also suggest, rather surprisingly, that even with numeral restrictors such as *Each of these five circles...*, as indicated in Figure 2, participants can tolerate a proper subset satisfying the presupposition in other visual arrangements.)




Figure 2: Illustration of item used in Geurt and van Tiel’s (2016) Experiment 1

In conclusion, DR forms an important factor that affects which inferences result from sentences involving projection from a quantificational context. This constitutes an important possible confound in assessing whether a given sentence gives rise to existential or universal projection, since a seemingly non-universal inference may in fact be the result of DR. Our experimental design is set up to allow for a differentiation of genuine existential projection from universal projection weakened by DR.

2.2. Entailed vs. non-entailed presuppositions

A second confound in distinguishing an existential from a seemingly non-universal inference is introduced if we allow for the possibility that presupposed content is also part of the conventionally entailed content (Sudo, 2012; Klinedinst, 2016). Let us illustrate with (6), which presupposes that the alien was blue at an earlier stage. However, with respect to the assertive meaning of this sentence, things are less clear. If the presuppositional content is simultaneously part of the entailed content, the conventional entailment would be as in (6a). In contrast, if the presupposition is not part of the entailed content, the conventional entailment is as in (6b).

- (6) On Planet B, the alien turned blue *again*. CONVENTIONAL ENTAILMENT:
- a. **Presupposition also entailed:** the alien turned blue & *was blue at an earlier stage*
 - b. **Presupposition entirely separate:** the alien turned blue

Note that, while (6b) and (6a) are contextually equivalent given the presupposition, the contrast between entailment versus non-entailment of the presupposition could interact with projection. In particular, universal quantifiers yielding a universal *reading* for a presupposition that is also entailed would simply be the result of assessing the quantificational claim relative to the entailed content (which, by hypothesis, includes the presupposition) and thus does not necessarily indicate universal *projection*. A key question is how we settle the issue of whether a given trigger simultaneously introduces its presupposition as an entailment as well. Indeed, Sudo (2012) argues that presupposition triggers can differ precisely in this regard, and Zehr and Schwarz (2016) provide some initial evidence from non-monotonic quantifiers (where the predictions come apart most clearly, as observed by Sudo). Our working hypothesis, building on prior work, is that *aware* is a good candidate for entailing its presupposition (Djäv et al., 2017), while *again* does not seem to entail its presupposition (Zehr and Schwarz, 2016).

A further complication arises when considering the possibility of presupposed content forming part of the entailments as well is that most theories of presupposition allow for some version of local accommodation (Heim, 1983), which effectively turns presupposed content into entailed content, while cancelling its contribution qua presupposition. This is distinct from the notion of entailed presuppositions we just introduced, but not easy to tease apart empirically. Different types of triggers are commonly thought to differ in how easily available local accommodation is, with triggers like *again* showing more resistance to such readings. See Klinedinst (2016) for a discussion of local accommodation vs. entailed presuppositions.

2.3. The design of the present study

Putting the various factors together (universal vs. existential projection, DR, and (non-)entailment), there is a total of five different logically possible readings for the presupposed content in a given quantifier-trigger combination. The first possible reading is a universal reading that is derived directly from universal presupposition projection (\forall in Table 1) and that yields an unrestricted universal inference throughout. The second reading is an existential reading that is derived directly from existential presupposition projection and in which the presupposition features in the entailed content (\exists + EntPS), and which yields universal inferences for universal quantifiers only. The third reading is an apparent existential reading (relative to the full domain) that is derived from universal projection by DR. The fourth reading is an existential reading that is derived from existential presupposition projection and in which the presupposition does not feature in the entailed content (\exists + no EntPS). And finally, the fifth reading is a presupposition-less reading (PS-LESS in Table 1), as presuppositions are well-known to be subject to suspension or cancellation. These readings are illustrated in Table 1 with the factive trigger *aware* (left) and the iterative trigger *again* (right).²

We designed an experiment to further investigate the projection behavior of presuppositions in the scope of quantifiers, and, specifically, to tease apart the different readings in Table 1 within a single design. The question we aim to answer is whether, in case of an existential

²The schematic pattern used in the table is shorthand for (based on the example with *aware* in row 1) ‘all aliens in fact ARE blue, and Q aliens THINK they are blue’, where Q is the relevant quantifier.

	<i>Q</i> alien is aware that he is blue	<i>Q</i> alien turned blue again
1. \forall	all are, <i>Q</i> thinks blue	all were, <i>Q</i> is blue
2. \exists +Entailment	one ⁺ is, <i>Q</i> is & thinks blue	one ⁺ was, <i>Q</i> was & is blue
3. DR	<i>Q</i> blue thinks blue	<i>Q</i> previously blue is blue
4. \exists −Entailment	one ⁺ is, <i>Q</i> thinks blue	one ⁺ was, <i>Q</i> is blue
5. PS-LESS (−Ent.)	<i>Q</i> thinks blue	<i>Q</i> is blue

Table 1: The five different possible readings for the interpretation of presuppositions in the scope of quantifiers, in which *Q* stands for *every* or *at least one*.

presupposition, the presupposition is derived directly via presupposition projection or via the predicted presupposition in combination with some other mechanism (DR, (non-)entailment). Moreover, we use two different triggers (*aware*, *again*) and two quantifiers (*every*, *at least one*) to test whether projection behavior differs across triggers and/or quantifiers.

3. Methods

3.1. Materials & Design

We use a picture-matching task with a (partially) covered box (Huang et al. 2013). In a covered box task, subjects are asked to select a match for a given sentence among various pictures, one of which is hidden. The covered box allows for a choice that better fits with subjects' expectations without making it salient, thereby avoiding a situation in which they must give either a direct *yes* or *no* response when neither seems quite appropriate, as is often the case due to presuppositional requirements.

The experiment consists of two sub-experiments: one with the factive trigger *aware*, and one with the trigger *again*. Each sub-experiment includes the quantifiers *every* and *at least one*. For the sub-experiment with *aware*, sentences of the form *Q* alien is aware that he is color were used (in which *Q* stands for either *every* or *at least one*, and the specific color differed per trial). To establish a plausible context, participants were told that the aliens cannot directly perceive their own skin color, and that they can only find out what color they have through the use of a machine, which sometimes may malfunction, leading to wrong ideas about their own color. Written sentences were presented along with two pictures of seven aliens (see Figure 3). The aliens' actual color represents the presuppositional dimension and the thought bubble-renderings of the aliens' beliefs represent the assertive dimension. In the 'covered box' picture, the aliens and thought bubbles were hidden by black squares.




Figure 3: Example item for a sentence like *Every/at least one alien is aware that he is blue*.

In the sub-experiment with the trigger *again*, sentences like *Q* alien turned color again were

used. These sentences were paired with pictures showing aliens traveling from a home planet on which they had a certain color (the presuppositional dimension) to a planet that we called Planet A on which they all lost their color, indicated by showing them as gray, and finally to a third planet (Planet B) on which they turned a color (other than gray) again. Planet B represents the assertive dimension. In the covered box picture, the aliens on the home planet and on Planet B are covered with black boxes. This is shown in Figure 4.

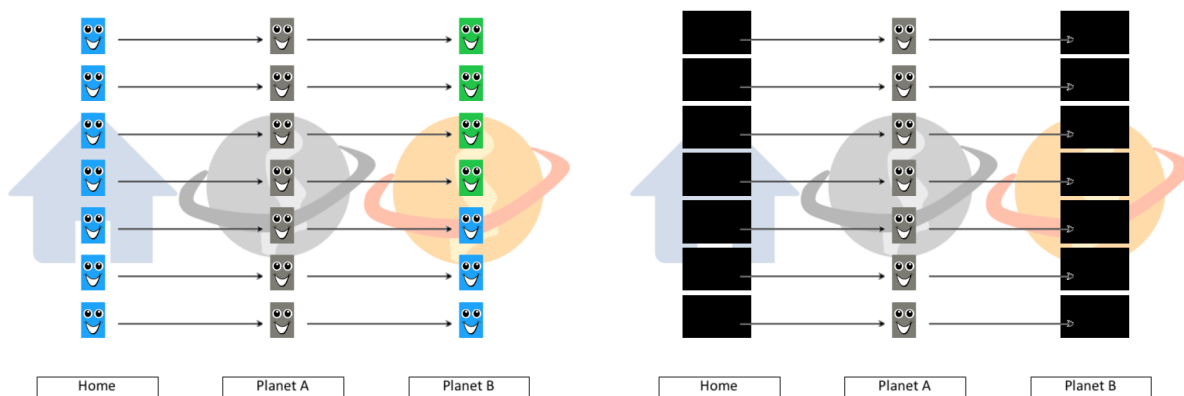


Figure 4: Example item for a sentence like *Every/at least one alien turned blue again*.

We included 6 conditions for each quantifier, of which 3 were critical conditions (\exists PS1, \exists PS2, and FALSEPS) and 3 were control conditions (FALSEASSERTION, ALLTRUE, and ALLFALSE).

In addition to the types of materials introduced in detail, a block of sentences with the negative quantifier *no* was included, both to see whether it exhibited a pattern closer to *at least one* or *every* and also whether the effects for *every* and *at least one* would be affected by seeing the block of sentences with *no* before or after the block that included *every* or *at least one*. However, the results for *no* were complex in a way that goes beyond what we have space for in the present paper, and as there was no significant impact of *no*-blocks preceding either one of the other quantifiers, we will collapse the data for those quantifiers from subjects seeing different block orders, yielding one group of subjects that saw trials with *at least one* (either before or after a *no*-block) and another that saw trials with *every* (again, in either order with the *no*-block). There were 5 items per condition per quantifier, so that every subject saw 30 items with *at least one* or *every* and 30 items with *no*, counterbalanced so that each item was only seen in one condition. The next section lays out the conditions that were used in the experiment in detail.

3.2. Conditions & Predictions

The conditions consist of picture-variations using different color distributions, which in turn yield varying compatibility with the candidate interpretations as defined in Table 1 above. The color distributions are equivalent in the two sub-experiments with *aware* and *again*: rotating the pictures for *again* 90° counter-clockwise shows the similarity with the pictures for *aware*. Each condition displays different pictures for the quantifiers *at least one* and *every*, to account for the interplay of quantifier and the various factors affecting the resulting presupposition reading.

The predictions in terms of compatibility with the candidate interpretations that were defined in Table 1 are given in Table 2 for *at least one* and in Table 3 for *every*.

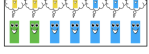


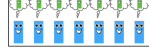


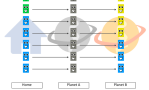





Aware							
Again							
PS reading		\exists PS1	\exists PS2	FALSEPS	FALSEASS.	ALLTRUE	ALLFALSE
\forall		✗	✗	✗	✗	✓	✗
\exists +ENT.		✓	✗	✗	✗	✓	✗
DR		✓	✗	✗	✗	✓	✗
\exists -ENT.		✓	✓	✗	✗	✓	✗
PS-LESS		✓	✓	✓	✗	✓	✗

Table 2: Predictions for the quantifier *at least one* in a sentence like **At least one** alien {is *aware* that he is blue / turned blue *again*} in the 6 conditions for the triggers *aware* (top row of images) and *again* (bottom row of images).

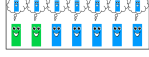
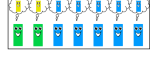
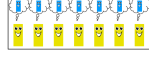
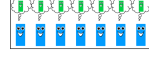
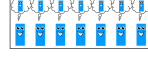
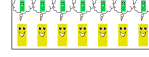
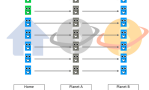





Aware							
Again							
PS reading		\exists PS1	\exists PS2	FALSEPS	FALSEASS.	ALLTRUE	ALLFALSE
\forall		✗	✗	✗	✗	✓	✗
\exists +ENT.		✗	✗	✗	✗	✓	✗
DR		✓	✓	✗	✗	✓	✗
\exists -ENT.		✓	✗	✗	✗	✓	✗
PS-LESS		✓	✗	✓	✗	✓	✗

Table 3: Predictions for the quantifier *every*, in sentence like **Every** alien {is *aware* that he is blue / turned blue *again*} in the 6 conditions for the triggers *aware* (top row of images) and *again* (bottom row of images).

The critical conditions for *at least one* are set up as follows: (i) \exists PS1 is incompatible with a universal reading, since there are aliens that are not blue, but it is compatible with all other readings. (ii) \exists PS2 is only compatible with two readings: an existential reading where the presupposition is entirely separate from the entailed content, as there is no alien that both is blue and thinks that they are blue, and a presupposition-less reading, which merely requires there to be at least one alien that thinks that they are blue. (iii) FALSEPS is only compatible with a presupposition-less reading, since there are no aliens that actually are blue. The comparison between the latter two thus will be indicative of the existence of an existential reading without an entailed presupposition (and there is no independent requirement for the presupposition and

entailed content to hold of the same individual(s); cf. the ‘Binding problem’ for presuppositions), as higher levels of acceptance in \exists PS2 could only be due to the availability of such a reading. The comparison between \exists PS1 and \exists PS2 will indicate to what extent either DR or an entailed presupposition is at play.

The pattern of compatibility with the conditions for *every* is slightly different. (i) In addition to being incompatible with the universal reading, as for *at least one*, \exists PS1 is incompatible with existential projection and an entailed presupposition, since *every* requires all entailed content to hold universally. (ii) \exists PS2 can only be accepted under a DR reading, since not all aliens are blue, and not all aliens think that they are blue. Only if the sentence is evaluated relative to a domain restricted to blue aliens can it be accepted. (iii) As before, FALSEPS is laid out so that the overt picture can only be accepted under a presupposition-less reading, as there are no blue aliens, but all aliens think they are blue.

Control items implement the same conceptual manipulation with adjustments as necessary for whichever quantifier is used: (iv) For FALSEASSERTION, the assertion is false since none of the aliens think that they are blue, but the presupposition is universally met, since all aliens actually are blue. (v) ALLTRUE completely fits with both universal presupposition and the respective asserted requirements (regardless of entailment). Finally, (vi) for ALLFALSE, neither the presuppositional requirements (on any variant) nor the assertive ones are met. These control items serve to assess participants’ understanding of the task and provide points of reference at both the ceiling and floor levels.

3.3. Participants & Procedure

160 undergraduate students at the University of Pennsylvania took part in the experiment for course credit. Half of them took part in the sub-experiment with *aware*, and the other half with *again*. Both sub-experiments varied the quantifiers *every* and *at least one* as a between-subjects factor. The experiment was implemented in Ibex. The presentation order and whether the covered box appeared on the left or on the right was randomized in Ibex.

Participants were seated in front of a computer and were told that they have to determine which of two pictures corresponds to a sentence description. The experiment started with instructions that showed participants the aliens they would be seeing. For *aware*, participants were told that the aliens are not able to directly perceive their skin color. Rather, a sometimes dysfunctional machine informs the aliens about their color. For *again*, participants were told that the aliens change color going from planet to planet. They travel from their home planet to planet A (where they turn gray), and on to planet B. Participants were instructed to press the F key on their keyboard to accept the left picture, and the J key to accept the picture on the right. We included a couple of practice trials with feedback, after which the actual experiment started.³

³Archive versions of the experiment can be found at <http://spellout.net/ibexexps/SchwarzLabArchive/AvaQPsAgain/> (**Again**) and <http://spellout.net/ibexexps/SchwarzLabArchive/AvaQPsAware/> (**Aware**).

4. Data analysis & results

We used logistic regression mixed effect models to predict the choice of the visible picture, using the lme4 package (Bates et al., 2015: version 1.1 – 13) in the R environment (version 3.3.3). We computed models on each pair of conditions for which different readings make different predictions, for the quantifiers *at least one* and *every* (see Tables 2 and 3). In addition to Condition, all the models whose outputs we report here included another two-level predictor: Trigger (*aware* = -1 vs. *again* = 1). In our reports below, we always mention the condition coded as -1 first, and the condition coded a 1 second. The models tested both for simple effects and for interactions between the two factors. Following the procedures for model simplification in Bates et al. (2015), we fitted models with a maximal random-effect structure (random slopes for Condition per Subject and random intercepts for items) and proceeded to an iterative reduction. As a result, when appropriate, we report models that forced a zero correlation on Condition per subject (using the ‘||’ syntax of `lmer()`). Whenever both types of models would converge, their outputs were qualitatively equivalent. Goodness of fit was reported to significantly decrease in all models dropping the random slope for Condition per Subject.

Besides the unfiltered data set, we ran models with data sets that excluded subjects with an accuracy of under 65% (excluding 4 participants), under 70% (excluding 6 participants), and under 75% (excluding 11 participants) on the ALLTRUE and ALLFALSE items. Since this filtering on accuracy hardly ever made a difference in terms of eliciting significant contrasts, we report the outputs of the models run on untrimmed data sets, except when the models failed to converge in which case we report the next most conservative converging model.

The results are presented in Figure 5 for the quantifier *at least one* and in Figure 6 for *every*. We start by discussing the results for *at least one*, after which we discuss the results for *every*. In Section 4.3, we discuss the data in terms of different sub-groups, which indicates that individual participants (consistently) adopted different strategies.

4.1. Results *at least one*

The results for *at least one* are presented in Figure 5. As expected, target acceptance rates for the ALLFALSE, FALSEPS, and ALLTRUE conditions are at floor and ceiling, respectively, for both triggers. Note further that the results on the different conditions are very similar for the different triggers. In our analysis, we first compared the choice of visible picture in the conditions \exists PS1 and ALLTRUE. A contrast here would be indicative of unrestricted universal projection (\forall) (see the \forall row in Table 2). Neither the full model nor the zero-correlation model converged. The next most conservative converging and parsimonious model is a zero-correlation model with the data set that is filtered for 65% accuracy. The model does not reveal a significant difference ($\beta = 0.0811, SE = 0.2620, p = 0.7571$), indicating that unrestricted universal presupposition readings were not at play for our participants in responding to the *at least one* items. Second, we compared \exists PS1 and \exists PS2. Both of these should be accepted across the board if (‘unbound’) \exists -ENTAILED or PS-LESS readings are widely available. A contrast between the two would point to \exists +ENTAILED or DR




Figure 5: Results for the items with the quantifier *at least one*, for sentences like: **At least one** alien {is aware that he is blue / turned blue again} with the triggers *aware* and *again*.

readings being available. The model shows a significant difference between \exists PS1 and \exists PS2 ($\beta = -9.7324, SE = 1.1360, p < 0.001$). This suggests that \exists -ENTAILED and PS-LESS readings are at most available to a limited extent, while \exists +ENTAILED or DR are responsible for ceiling level acceptance of the target in condition \exists PS1. Third, we compared the choice of visible picture in \exists PS2 and FALSEPS. A contrast between these conditions would point to the availability of \exists -ENTAILED readings (which allow acceptance of \exists PS2). The model shows a significant contrast ($\beta = -0.5742, SE = 0.2828, p = 0.0423$), indicating that \exists -ENTAILED readings might have been at play for our participants, though apparently only to a limited extent, given the small size of the effect, which also is numerically more pronounced for *again* (though note that there is no significant interaction between triggers). Fourth, we compared the choice of visible picture in \exists PS2 and ALLFALSE. Here, both \exists -ENTAILED and PS-LESS readings predict a difference between the two conditions (acceptance for \exists PS2, rejection for ALLFALSE). While the numeric difference is quite large, the model does not reveal such a contrast ($\beta = -0.1937, SE = 0.2870, p = 0.4998$), thereby not providing direct evidence for a \exists -ENTAILED or a PS-LESS reading (see below for discussion of relevant differences in individual subjects' response patterns). Finally, we compared FALSEPS and ALLFALSE, for which only a PS-LESS reading predicts a contrast. As with \exists PS2, the model does not reveal a significant contrast ($\beta = -0.1131, SE = 0.3652, p = 0.7569$), despite a fairly large numerical difference (*again*, see discussion of individual differences below).

To summarize, the only significant differences our models detected were between the \exists PS1 and \exists PS2 conditions and between the \exists PS2 and FALSEPS conditions. The contrast between \exists PS1 and \exists PS2 suggests that for both triggers, either a reading where the presupposition is also part of the entailed content is available, or else one based on DR (though this seems less likely, given the results for *every* below). As was noted above, presuppositions can wind up contributing to entailed content directly in at least two ways, as local accommodation can render a comparable

result (while removing the presuppositional component altogether), and our results here do not differentiate between these possibilities. The contrast between \exists PS2 and FALSEPS suggests that, to a limited extent, subjects accepted \exists PS2 under a reading where the presupposition is not part of nor bound to the entailed content. Numerically, the difference seems to be bigger for *again* than for *aware*. We will discuss the availability of \exists -ENTAILED and \exists +ENTAILED readings in more detail in Section 4.3, after discussing the results for *every* in the next section.

4.2. Results *every*

The results for *every* are presented in Figure 6. Again, the target acceptance rates for the ALLFALSE, FALSEPS, and ALLTRUE conditions are at floor and ceiling, respectively, for both Triggers. Similar to the data for *at least one*, the results on the different conditions pattern similarly for the two triggers. However, for \exists PS1, the results are quite different from those for *at least one*, pointing to clear differences in descriptive projection patterns between quantifiers.




Figure 6: Results for the items with the quantifier *every*, for sentences like: **Every** alien {is aware that he is blue / turned blue again} with the triggers *aware* and *again*.

The first model we ran compared the \exists PS1 and ALLTRUE conditions, for which generally available unrestricted \forall projection and \exists +ENTAILED both predict a difference (with either predicting rejection of \exists PS1). Indeed, our model shows a significant contrast between the two conditions ($\beta = 1.0712, SE = 0.2402, p < 0.001$). Second, we compared \exists PS1 and ALLFALSE, for which DR, \exists -ENTAILED, and PS-LESS readings predict a difference (acceptance for \exists PS1). The model reveals a significant contrast between the conditions ($\beta = -1.0887, SE = 0.2164, p < 0.001$). Third, we compared \exists PS1 and \exists PS2. If there is no significant contrast between these conditions, this would show that Domain Restriction accounts for all of the acceptances in \exists PS1. However, the model shows a significant contrast ($\beta = -4.3254, SE = 0.7970, p < 0.001$), which suggests that DR cannot account for all of the difference between \exists PS1 and \exists PS2. Rather, some of the acceptances of \exists PS1 must be based

on a \exists -ENTAILED or a PS-LESS reading. Fourth, we compared \exists PS1 and FALSEPS to assess the extent to which acceptance in the former is driven by a PS-LESS reading. We find a significant contrast ($\beta = -0.7489$, $SE = 0.1899$, $p < .001$), suggesting that not all such responses are based on this reading. Fifth, we compared the \exists PS2 and ALLFALSE conditions, for which only a reading that follows from DR predicts a difference. No such contrast was revealed by our model ($\beta = -0.1761$, $SE = 0.2788$, $p = 0.5276$), suggesting that DR does not play a role. However, in the next section we will discuss the individual results, which show that, even though DR might be limited, there are some subjects with high acceptance rates for \exists PS2. Finally, we compared the acceptances of the overt picture for FALSEPS and ALLFALSE, for which only a PS-LESS reading predicts a difference. Again, the model did not reveal a significant contrast ($\beta = -0.2736$, $SE = 0.3923$, $p = 0.4855$). However, again, even though the model does not show a significant contrast, there appears to be a proportion of the subjects with high acceptance rates for FALSEPS. We will discuss this in the next section.

To summarize, the results for *every* show a significant difference between \exists PS1 and ALLTRUE and between \exists PS1 and ALLFALSE. The first finding provides clear evidence that presuppositions triggered from the scope of *every* have a universal projection (and/or a \exists +ENTAILED projection; see discussion below). The second contrast could be driven by a DR reading, a \exists -ENTAILED reading, or a PS-LESS reading. Clearly, these results require closer inspection. Several additional aspects of the data indicate that a closer look at the distribution of the answers of the different participants is in order. First, the results on most critical conditions (in contrast to the control conditions) for both *at least one* and *every* do not show 100% acceptances or rejections. This could be caused by a bimodal distribution in the acceptance rates (inter-subject differences). Moreover, the comparison between \exists PS1 and \exists PS2 and between FALSEPS and ALLFALSE does not provide a significant difference, although we see quite large numerical differences as well as some individuals who have high acceptances on \exists PS2 and on FALSEPS. Finally, note that we found a much greater fit for models allowing for random slopes for Condition per Subject. This shows that the slopes capture significant variation in effect size per subject; the models with random slopes therefore reduce the residual variance. However, the models that drop the random slopes (which we fit to arrive at the most parsimonious model in terms of the random-effects structure) indicate significant differences where conditions descriptively appear to contrast. This is for instance the case with the two final comparisons that we discussed in this section: between \exists PS2 and ALLFALSE and between FALSEPS and ALLFALSE. We need to further inspect the data to see whether the variance that is captured by the slopes actually comes from a “real” difference that is masked by the random slope. In the next section, we therefore explore the individual participants’ profiles, and we show that there are different speaker populations.

4.3. Participants’ profiles

Based on the results, this section further inspects the different participants’ profiles, first for *at least one*, and then for *every*.

4.3.1. At least one

While statistical comparisons of \exists PS2 and ALLFALSE and of FALSEPS and ALLFALSE did not reveal significant contrasts, we found fairly large numerical differences. We inspected individual participants' profiles on the FALSEPS and \exists PS2 conditions to better understand the source of the numerical contrasts, in particular with regards to participants' consistency in responses as well as the individual response patterns across relevant conditions. Recall that FALSEPS can only be accepted under a PS-LESS reading, and that \exists PS2 can be accepted under a PS-LESS reading as well as a \exists -ENTAILED reading. Figure 7 plots the mean acceptance on FALSEPS on the y-axis and the mean acceptance on \exists PS1 on the x-axis. While most participants reject the overt picture in both conditions, there is a small but not insubstantial number of subjects who consistently accept the overt picture, especially for *aware*. Furthermore, the roughly linear increase in the distribution suggests a correlation between accepting FALSEPS and \exists PS2, which is expected if acceptance is based on their ability to access a PS-LESS reading. There are a few subjects that diverge from this distribution. This is clearer for *again* than for *aware*: these subjects consistently accept the overt picture in \exists PS2 but not in FALSEPS, suggesting that they access a \exists -ENTAILED reading.




Figure 7: By-subject plot on the conditions \exists PS2 and FALSEPS with the quantifier *at least one* and the triggers *aware* (left) and *again* (right).

4.3.2. Every

For *every*, our models similarly failed to find a significant contrast between ALLFALSE and \exists PS2 on the one hand and FALSEPS on the other, despite non-negligible-seeming numerical contrasts. The former contrast would indicate the availability of DR, and the latter a PS-LESS reading. We inspected individual response distributions to assess the source of the sizable numerical effects as well as potential individual response patterns.

Starting with FALSEPS, the y-axis distribution in Figure 8 shows that a large proportion of

subjects—close to half—accepts the target in this condition quite consistently, often at ceiling levels. This suggests that PS-LESS readings are systematically available for some subjects. While the significant contrast with \exists PS1 reported above suggests that not all acceptances in that condition are due to PS-LESS readings, it is still possible that a large portion of them are. Indeed, comparing the x-axis to the \exists PS1-response distribution on the y-axis in Figure 8 suggests a strong correlation between these two conditions for *aware*, as participants who accept \exists PS1 also accept FALSEPS, while participants who reject \exists PS1 also reject FALSEPS, with only a couple of exceptions. This indicates that a PS-LESS reading of *aware* is consistently available for at least some of the participants, and furthermore that to a large extent, acceptance in \exists PS1 is also driven by this reading (rather than DR or an \exists -ENTAILED reading).




Figure 8: By-subject plot on the conditions \exists PS1 and FALSEPS with the quantifier *every* and the triggers *aware* (left) and *again* (right).

Interestingly, it is less clear that there exists such a correlation for *every* with the trigger *again*. As can be seen on the right side in Figure 8, with *again*, there are some participants who systematically reject the overt picture for FALSEPS while still accepting \exists PS1 in at least a portion of the cases. Note that \exists PS1 could be accepted under a reading that comes from DR, from a \exists -ENTAILED reading, or from a PS-LESS reading. It is safe to conclude that those people who reject FALSEPS do not accept \exists PS1 under a PS-LESS reading. Note further that \exists PS2 can only be accepted under a DR reading, and that we see a much higher acceptance of \exists PS1 compared to \exists PS2 in Figure 6. Therefore, it is likely that an explanation for the difference between the presupposition triggers should be sought in the idea that triggers can differ with regards to whether or not their presupposition also features in the entailed content (Sudo, 2012). In line with results from previous studies (Djäv et al., 2017; Zehr and Schwarz, 2016), our results suggest that *aware* entails its presupposition (unless one has a PS-LESS interpretation of *aware*), while *again* does not entail its presupposition.

Finally, Figure 9 plots the distribution of results across participants for \exists PS2, to further investigate the role of DR, which our overall statistical results suggest is quite limited. The histogram shows that, while the overwhelming majority of the subjects never accept the overt picture for \exists PS2, there are some subjects who (sometimes) accept the picture and, thus, necessarily apply

DR. In total, there are 4 subjects who have high acceptance rates ($>80\%$) for \exists PS2 with *aware* and 7 subjects with *again*. This indicates that, although very limited, DR is an available reading for some of the subjects in our study.




Figure 9: Histogram of the mean acceptance rate on \exists PS2 with the quantifier *every* and the triggers *aware* (left) and *again* (right).

5. Discussion

The experiments in this paper were set up to address a) whether presupposed content under quantifiers gives rise to universal or existential presupposition-based inferences and b) whether the observed reading(s) arise(s) as a direct result of the projection mechanism or via additional mechanisms such as Domain Restriction and (non-)entailment of the presuppositional content. We used two different triggers (*aware*, *again*) and two quantifiers (*every*, *at least one*) to test whether projection behavior differs across triggers and/or quantifiers.

Our data provide clear confirmation that presupposition-based inference patterns vary by quantifier, in line with previous results by Chemla (2009) and Tiemann (2014): targets with non-universally met presuppositions are readily accepted for *at least one* but much less often so for *every*. Moreover, the results show that the overall results pattern is quite comparable across the two triggers, although we do observe subtle differences in terms of entailment for the different triggers.

As for determining how the descriptively universal and existential readings for the two quantifiers should be accounted for in theoretical terms, a detailed consideration of the various factors at play is required. The rejections of \exists PS1 for *every*, reflecting a descriptively universal inference, can be accounted for either in terms of universal projection (\forall) or via existential projection plus an entailed presupposition. However, if we assume that *again* does not entail its presupposition, as suggested by prior work as well as by some aspects of our results, then this finding is indeed supportive of \forall -projection from the scope of *every*.

Importantly, the extent to which responses indicating non-universal readings for *every* can be attributed to DR in our data seems to be extremely limited, as we find no general statistically significant effects directly attributable to DR. This is in contrast to previous findings by Geurts

and van Tiel (2016). At the same time, however, there are some individual participants who show consistent acceptance of targets that are only compatible with universal projection relative to a restricted domain, suggesting that this interpretative option is in principle available but only accessible to few speakers in our experimental context.

In addition to the variation in presupposition-based inference patterns between quantifiers, we also find variation between readings that incorporate presupposed content as part of the entailed content and readings that do not. This is the case for both triggers, although there are some indications of differences between triggers as well. For the quantifier *at least one*, we find that a fair number of participants accept overt pictures that are only compatible with a \exists -ENTAILED reading or a PS-LESS reading. The former would be expected, to some extent, for *again*, based on results from prior work arguing it to be a non-entailing trigger. For *aware*, the availability of either reading is more surprising. However, it is quite plausible that in this case, the result is attributable to the specific nature of the task at hand rather than a lexical property of *aware*. In particular, given the context provided in the instructions, where the aliens rely on sometimes faulty machines to form beliefs about their own color, there may be a notion of *aware* that takes into account a perspectival shift of sorts: as far as the alien in question is concerned, they may perfectly well have reasonably justified belief about their color based on the machine-feedback, even if that feedback could be faulty, as that is the only source of information at their disposal. It is in light of this justification from the perspective of the alien that one could describe them as ‘aware’ of their color, even if they wind up getting the color wrong.⁴ What appear to be ‘presupposition-less’ readings in descriptive terms might then be regular presuppositional readings with some shift in perspective.

Turning to *again*, we find some evidence for both types of readings, but these may need to be accounted for in different terms. First, if we assume (following previous work) that *again* does not entail its presupposition, the observation of \exists -ENTAILED is straightforwardly accounted for.⁵ A non-entailed representation of *again* also accounts for the responses of participants who, at the same time, rejected overt pictures only compatible with PS-LESS readings, and leaves open the possibility of local accommodation to account for the responses of the participants who showed no evidence of accessing \exists -ENTAILED or PS-LESS readings of *again*. However, it is in principle possible that in certain circumstances, the presupposition of *again* can simply be ignored, which accounts for observations of PS-LESS readings.

6. Conclusions

Presuppositions give rise to different inference patterns from different quantifiers, as documented here for universal and existential ones. Theoretical accounts of these differences are complicated by a variety of factors, such as Domain Restriction and (non-)entailment of presuppositions. In light of previously proposed differences between types of triggers, our results suggest that the projection mechanism itself yields universal and existential readings from the respective quantifiers and that Domain Restriction at best plays a very limited role. At the same

⁴Thanks to Jeff Lidz for first spelling out this possibility for us in fully explicit terms.

⁵Though note that Sudo (2012) proposes cross-dimensional anaphora to account for *exactly one* binding its quantified variable in the presuppositional as well as in the assertive dimension, thus predicting rejection in \exists -ENTAILED, where no alien satisfies *both* dimensions at the same time.

time, there is substantial variation in the types of readings that are possible for these presupposition triggers, and, ultimately, further work is needed to pin down which theoretical properties the various interpretative effects should be attributed to. This will also require the investigation of a wider range of triggers and quantifiers.

References

- Bates, D., R. Kliegl, S. Vasishth, and H. Baayen (2015). Parsimonious mixed models. *arXiv preprint arXiv:1506.04967*.
- Bates, D., M. Mächler, B. Bolker, and S. Walker (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67(1), 1–48.
- Beaver, D. (1994). When variables don't vary enough. In *Semantics and linguistic theory*, Volume 4, pp. 35–60.
- Beaver, D. I. (2001). *Presupposition and assertion in dynamic semantics*, Volume 16. CSLI publications Stanford.
- Chemla, E. (2009). Presuppositions of quantified sentences: experimental data. *Natural Language Semantics* 17(4), 299–340.
- Djäv, K., J. Zehr, and F. Schwarz (2017). Cognitive vs. emotive factives: An experimental differentiation. *Proceedings of Sinn und Bedeutung* 21.
- Fox, D. (2012). Presupposition projection from quantificational sentences: trivalence, local accommodation, and presupposition strengthening. MS the Hebrew University of Jerusalem.
- Geurts, B. (1999). *Presupposition and pronouns*. Elsevier, Oxford.
- Geurts, B. and B. van Tiel (2016). When all the five circles are four: New exercises in domain restriction. *Topoi* 35(1), 109–122.
- Heim, I. (1983). On the projection problem for presuppositions. *Formal semantics—the essential readings*, 249–260.
- Karttunen, L. (1973). Presuppositions of compound sentences. *Linguistic Inquiry* 4(2), 169–193.
- Klinedinst, N. (2016). *Two Types of Semantic Presuppositions*, pp. 601–624. Cham: Springer International Publishing.
- Rothschild, D. (2011). Explaining presupposition projection with dynamic semantics. *Semantics and Pragmatics* 4, 3–1.
- Schlenker, P. (2008). Be articulate: A pragmatic theory of presupposition projection. *Theoretical Linguistics* 34(3), 157–212.
- Schlenker, P. (2009). Local contexts. *Semantics and pragmatics* 2, 3–1.
- Sudo, Y. (2012). *On the semantics of phi features on pronouns*. Ph. D. thesis, Massachusetts Institute of Technology.
- Tiemann, S. (2014). *The processing of 'wieder' ('again') and other presupposition triggers*. Ph. D. thesis, Universität Tübingen.
- Zehr, J., C. Bill, L. Tieu, J. Romol, and F. Schwarz (2016). Presupposition projection from the scope of *None*: Universal, existential, or both? In *Proceedings of SALT* 26.
- Zehr, J. and F. Schwarz (2016). Entailed vs. non-entailed presuppositions - an experimental assessment. In *Proceedings of NELS 46*, Amherst, MA. GLSA Publications.